
DOES GOVERNMENT DECENTRALIZATION INCREASE POLICY INNOVATION?

KOLEMAN S. STRUMPF

University of North Carolina at Chapel Hill

Abstract

The conventional wisdom is that government decentralization promotes policy innovation because it allows for several simultaneous experiments by local governments. However, this ignores a learning externality: successful policy experiments provide useful information for all governments. Local governments will ignore this externality, but a central government should take it into account. This article uses a social learning model to compare policy innovation under centralization and decentralization. Centralization leads to more policy innovation if the local governments are relatively homogeneous or large in number. However, decentralization may induce more policy innovation if there are multiple experimental policies available.

It is one of the happy incidents of the federal system that a single courageous State may, if its citizens choose, serve as a laboratory; and try novel social and economic experiments without risk to the rest of the country.

Justice Louis D. Brandeis, *New State Ice Co. v. Liebmann*, 1932.

1. Introduction

One of the textbook arguments for government decentralization is that it promotes policy innovation.¹ As the above quote from Justice Brandeis

Koleman S. Strumpf, Department of Economics, University of North Carolina at Chapel Hill, Gardner Hall, CB# 3305, Chapel Hill, NC 27599-3305 (cigar@unc.edu).

I wish to thank participants at the 2000 Stanford Institute for Theoretical Economics (SITE) conference, the 1997 Wallis Institute of Political Economy conference (University of Rochester), the 1997 Public Choice Society meetings, the UNC-Duke theory seminar as well as two anonymous referees for their comments. Financial support from the Latané Fund for Interdisciplinary Research in the Behavioral Sciences is greatly appreciated.

Received March 2000; Accepted May 2001.

¹For example, see Rosen (1999), a standard undergraduate public finance textbook.

suggests, local governments such as states and municipalities are often thought of as “laboratories of democracy” with each one independently pursuing different policy experiments. Decentralizing policy choice to local governments has the advantage that several different policies can be considered simultaneously. In contrast the central government can examine only one policy at a time and so will more slowly uncover superior new policy choices. This argument has played a central role in the trend towards decentralization in the U.S. (*The New York Times*, 24 February 1997), with the most prominent example being the devolution of the welfare system to the states.

However, this reasoning ignores an important aspect of the policy innovation process. Because successful policy experiments are eventually emulated, they have a public good component. Experiments benefit not just the innovating government but also potential imitators, and so local governments have an incentive to free-ride off their neighbors. Alternatively, a central government should take this learning externality into account when it is deciding whether to consider a policy experiment.

This paper characterizes conditions under which decentralization (local government decision making) leads to more policy innovation than centralization (central government decision making). The model is a game theoretic version of the bandit problem. Policymakers must choose between a sure policy, one with a known payoff which can be thought of as the status quo, and an experimental policy, which may or may not have a superior payoff. For example, an experimental welfare policy involving time limits or worker training might greatly reduce future case-loads but could instead result in even higher poverty rates. Each policymaker must pick a sequence of policies with the objective of maximizing his present discounted flow of payoffs. The experimental policy may be the appropriate choice even when it has a lower expected current payoff because its use also reveals information about its underlying payoff distribution. Under decentralization, two local policymakers make decisions while observing one another’s decisions and payoffs. Although the experimental policy may be better suited to one policymaker than the other, a neighbor’s experience will provide valuable information because the local payoffs are correlated. The decentralized outcome typically involves under-experimentation relative to the social optimum because local policymakers do not take into account the informational externality their experiments provide. That is, a local policymaker may free-ride off his neighbor’s experiment.² Under centralization, a single policymaker must choose a uniform policy for both local governments. While the central policy-

²This point is related to other intergovernment externalities such as those from tax competition or spending spillovers. Here it is information rather than physical factors that flows between governments.

maker internalizes the learning externality, he only has access to a coarse set of policy instruments. The centralized outcome may involve under- or over-experimentation relative to the social optimum.

Centralization involves greater experimentation than decentralization if and only if there is a large positive correlation between the local experimental payoffs.³ This is because greater correlation increases the chance that the experimental policy is suitable for both local governments, and thus increases the appeal of centralized experimentation. Alternatively, greater correlation increases the severity of the free-rider effect and reduces decentralized experimentation. Increasing the number of local governments induces relatively more experimentation under centralization, since both the centralized experimental payoff and the free-rider effect grow with the number of governments. However, increasing the number of experimental policies may induce more decentralized experimentation. While centralization can only consider one experiment at a time, decentralization allows multiple policies to be used. Such decentralized policy diversity can be an equilibrium, since each local policymaker learns more from picking a unique experiment than from simply matching one of his neighbors. This provides a novel basis for Justice Brandeis's original conjecture: *free-riding* can promote decentralized experimentation. Finally, the model is generalized to allow for politically motivated local policymakers. While it is often argued that local policymakers introduce innovations as a means to get elected to higher political office, this result depends crucially on how the election process is modeled: political motivations may actually lead to less decentralized experimentation. In total these results suggest that contrary to conventional wisdom decentralization does not always induce more policy experimentation.

One crucial assumption in the model is that the central government is restricted to using a uniform policy in all regions. There are several reasons to think this is a realistic characterization. First, in many countries local governments have certain sovereign powers and cannot be forced to implement distinct policies from their neighbors. For example, the 10th Amendment to the U.S. Constitution is typically interpreted as reserving certain policies to the states. Second, there may be political resistance to differential policies emanating from the central government. Voters and representatives from regions where the least favorable experiments are to be used may try to block the policy. Third, even if the central government could simultaneously implement distinct policies it may be difficult for it to monitor or appropriately evaluate the results due to informational limits. Finally, in practice centralized policies are almost always uniform in industrialized countries.

³This result presumes the experimental policy is not expected to be extremely favorable or unfavorable for either local government.

This article is linked to two areas of active current research. The first involves the policy implications of federalist structures. The most closely related papers are Kollman et al. (1997) and Rose-Ackerman (1980). Kollman et al. (1997) theoretically investigate the role of central and local governments as innovators but impose rather than endogenize each government's learning strategy. Rose-Ackerman (1980) considers a model of local government innovation in which politicians are rewarded for their policy performance relative to their neighbors, but there is no strategic interaction between politicians. My paper builds on their work by allowing local governments to free-ride off one another's experiments. Allowing for such strategic interaction significantly alters the theoretical predictions. For example, Kollman et al. (1997) find that greater preference heterogeneity *reduces* the effectiveness of decentralization. Other papers consider how centralization or decentralization influences voting behavior. Besley and Coate (1999) show that centralization induces voters to select high-spending representatives to take advantage of the fiscal commons. Rose-Ackerman (1981) presents a model in which voter policy preferences differ at the central versus local level due to spillovers. My work is complementary to these papers because it focuses on strategic policymaking rather than voter behavior. Several recent papers consider the interaction between central and local governments. Persson and Tabellini (1996a, 1996b) examine how this interaction influences the riskiness of local policies, Dixit and Londregan (1998) consider its influence on redistribution policies, and Cremer and Palfrey (1998) examine the effect of federal mandates on local governments. Finally, the theory here can be connected to recent empirical work on federalism. Strumpf and Oberholzer-Gee (forthcoming) show that greater preference heterogeneity induces fiscal decentralization. This suggests that federal structures tend to be efficient, since I show that decentralization induces more policy experimentation when local governments are dissimilar. Very thorough surveys of the fiscal federalism literature are contained in Inman and Rubinfeld (1997) and Oates (1999).

The second research area involves learning externalities in multiple-person settings. Bolton and Harris (1999) analyze how several identical agents choose between a risky and safe action when they can observe one another. Bala and Goyal (1998) present a related model and an overview of the social learning literature. The contribution of my work is to allow for heterogeneous learning agents⁴ and more than two strategic actions.

⁴I am only aware of two papers in this literature that allow for inter-player experimental payoff heterogeneity. Smith and Sorensen (forthcoming) consider a sequential learning problem where each agent's type influences his payoff and is private knowledge. Bergemann and Valimaki (1997) have agents whose payoffs follow a known ranking. My model allows for a more general form of heterogeneity. The players' payoffs are correlated and their relative values may vary over time.

Also, the main focus of the social learning literature is determining whether the private equilibrium is socially optimal while my paper compares the private (decentralized) outcome with the constrained uniform (centralized) outcome. Comparing these second-best outcomes is potentially of interest in other settings. For example, an organizational theorist might be interested in which types of decisions are best decentralized within a firm (see Chang and Harrington, forthcoming, for a computational model of this problem).

2. Empirical Motivation

It is important to empirically document that governments copy policies from one another. Political scientists have extensively studied such emulation between U.S. states. The most well-known anecdote involves the California Fair Trade Law of 1931 which was copied verbatim by ten other states, including two serious typographical errors (Walker 1971). An earlier version of this paper contains many contemporary examples of states copying policies from their neighbors and delaying experiments in order to first evaluate other ongoing programs. There is also more formal empirical work which shows that policies diffuse between states.⁵ Berry (1994) and Nice (1994) contain detailed surveys of this literature.

This policy diffusion seems to be the result of active observation and contact between states. Freeman (1985) surveys state legislators and finds that they evaluate policies in other states before making their own policy proposals. Similarly, Walker (1971) concludes from his survey of state administrators that they frequently communicate with counterparts in other states. These contacts have developed into both formal and informal “policy networks” where information on new policies is exchanged (see the survey in Mintrom 1997). For example, the National Governors Association has established the Center for Best Practices which seeks to “identify and share states’ best practices and innovations” and serves as a clearing-house of policy evaluations (www.nga.org/CBP/About.asp).⁶ In addition to these politician networks, there are also several organizations that link professional support staff as well as organizations such as the Advisory

⁵These areas include judicial administration (Glick 1981), social services (Sigelman et al. 1981), energy policy (Regens 1980), computer use (Bingham 1976), tort law (Canon and Baum 1981), lottery adoption (Berry and Berry 1990), and the Progressive Movement (McCoy 1940).

⁶Similarly, The Council of State Governments’ Innovations Transfer Program “identifies and disseminates information on innovative programs and policies that have been successfully implemented by individual states and have the potential to be adapted for use in other states” (www.statesnews.org/publications/trends_innovations.html). The National Association of Counties Model County Program serves a similar role for local governments (www.naco.org/counties/models).

Commission on Intergovernmental Relations which provides model legislation based on state policy experiments.

While this brief review shows there is evidence of policy copying, the previous literature does not consider whether such emulation influences the rate of policy innovation. It is this connection that the remainder of the paper investigates.

3. The Model

3.1. Setup

This section lays out the basic model (extensions are considered in section 4). There are two regions and two periods. In each period either a sure policy, which has a known payoff, or an experimental policy, which has a payoff drawn from a known distribution, must be picked. The experimental payoffs in the regions may be different but have a known correlation. Under decentralization each region has a policymaker who selects a policy with the objective of maximizing the region's present discounted payoff. Under centralization a single central policymaker selects a common policy for both regions with the objective of maximizing the total present discounted payoff.

More formally, label the regions as $i = 1, 2$ and the periods as $t = 1, 2$. Let $a_{it} \in \{0, 1\}$ be the action choice in region i , period t . $a_{it} = 1$ indicates the experimental policy is used while $a_{it} = 0$ indicates the sure policy is used. Under decentralization, the strategy for each region's policymaker is an action for each period conditional on the history of play (mixed actions are allowed). Under centralization, the strategy for the central policymaker is an action to be applied in both regions for each period ($a_{1t} = a_{2t}$). The main focus will be on comparing the level of experimentation in each case. My measure of aggregate experimentation is the sum of the first period actions $a_{11} + a_{21}$. This experimentation measure excludes second period actions, since as will be made clear later all learning in this model is the result of first period actions.

Each policymaker discounts his second period payoff with a common factor, $\beta \in (0, 1)$. The payoff from the sure policy is normalized to zero in both regions. The payoff from the experimental policy is the sum of permanent and idiosyncratic components. The permanent component, θ_i , is fixed but is never observed. The idiosyncratic component, ϵ_{it} , changes each period and is also never observed.⁷ So, the payoff in region i , period t from the experimental policy is

$$\pi_{it} = \theta_i + \epsilon_{it}. \quad (1)$$

⁷The permanent term can be thought of as the expected net change in social welfare from implementing an innovative policy. The idiosyncratic term can be thought of as some fluctuating aspect of the economic environment which influences the policy outcome.

Notice that observing π_{it} does not reveal the exact value of θ_i because ϵ_{it} is also unknown. The permanent payoffs for the two regions, $\boldsymbol{\theta} \equiv \begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix}$, are drawn once from a bivariate normal distribution while the idiosyncratic payoffs are drawn each period from independent standard normal distributions

$$\boldsymbol{\theta} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \quad \epsilon_{it} \sim \mathcal{N}(0,1), \quad i, t = 1, 2. \quad (2)$$

The permanent payoff's mean vector, $\boldsymbol{\mu}$, and variance matrix, $\boldsymbol{\Sigma}$, may be written as

$$\boldsymbol{\mu} \equiv \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} \quad \boldsymbol{\Sigma} \equiv \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix}. \quad (3)$$

Here $\mu_i \in \mathfrak{R}$ is the mean and $\sigma_i^2 > 0$ is the variance in region i . Notice that these terms are allowed to vary across regions.⁸ $\rho \in [-1, 1]$ is the correlation between the experimental payoffs in the regions. The experimental payoffs are likely to be relatively similar when ρ approaches 1, relatively dissimilar when ρ approaches -1 , and independent when ρ is near 0. That is, $|\rho|$ measures the degree of dependence between the payoffs. All parameters, actions, and payoffs are presumed to be common knowledge.

There is an important asymmetry between the two periods. In the first period, the policymakers have beliefs about $\boldsymbol{\theta}$ based simply on the distribution in (2). In the second period, the policymakers may be able to refine these beliefs if the experimental policy was used in the first period. Notice that an experiment in one region provides useful information for the other region so long as $\rho \neq 0$.⁹ After the first period payoffs are observed, DeGroot (1970) shows that the Bayesian updated beliefs about $\boldsymbol{\theta}$ follow a normal distribution. The updated distribution's mean,

$\boldsymbol{\mu}^* \equiv \begin{pmatrix} \mu_1^* \\ \mu_2^* \end{pmatrix}$, will be important later. Prior to observing the first-period payoffs, the vector $\boldsymbol{\mu}^*$ is itself normally distributed with mean,

$$E(\boldsymbol{\mu}^*) = \boldsymbol{\mu} \quad (4)$$

and covariance matrix

$$\text{Cov}(\boldsymbol{\mu}^*) = \boldsymbol{\Sigma} - (\boldsymbol{\alpha} + \boldsymbol{\Sigma}^{-1})^{-1} \quad (5)$$

⁸This is because experimental policies that are successful in one region may be less suited to another due to differences in voter tastes, demographics, business composition, or economic health.

⁹The sure policy provides no information about the experimental payoff distribution.

where $\alpha \equiv \begin{pmatrix} a_{11} & 0 \\ 0 & a_{21} \end{pmatrix}$. The term in (4) represents the vector of expected experimental payoffs in the second period.

The diagonal terms on the covariance matrix in (5), which I will label as $\sigma_i^{*2}[a_{11}, a_{21}]$, play a crucial role in the analysis. $\sigma_i^{*2}[\bullet]$ measures how much μ_i is expected to move as a result of the first-period actions. When there are no experiments, the experimental mean remains unchanged and $\sigma_i^{*2}[0,0] = 0$. Alternatively, an experiment generates some information about the underlying payoff distribution and is expected to shift the mean. The amount of information revealed is increasing in the number of experiments. The following result shows that each $\sigma_i^{*2}[\bullet]$ is increasing in the level of aggregate experimentation.¹⁰

Remark 1 (Learning from Experimentation) For $i = 1, 2$,

$$\begin{aligned} \sigma_i^{*2}[1,1] &\geq \max(\sigma_i^{*2}[1,0], \sigma_i^{*2}[0,1]) \text{ and} \\ \min(\sigma_i^{*2}[1,0], \sigma_i^{*2}[0,1]) &\geq \sigma_i^{*2}[0,0] \end{aligned} \tag{6}$$

Proof: See Appendix A.1.

This discussion suggests $\sigma_i^{*2}[\bullet]$ should be interpreted as a measure of the amount of *learning* resulting from first-period experiments. As such, higher values of $\sigma_i^{*2}[\bullet]$ are beneficial (this will be proved in the next subsection).

3.2. Decentralization

In the *decentralized outcome* each policymaker selects actions that maximize his region’s expected present discounted payoff. The strategies are presumed to be subgame perfect, and so the equilibrium can be uncovered using backwards induction. In the second period policymakers seek to maximize their current payoff, since the game will end following this period. Each policymaker will select the policy that has the highest expected reward given his current information,

$$a_{i2} = \begin{cases} 1 & \mu_i^* \geq 0 \\ 0 & \mu_i^* < 0 \end{cases} \tag{7}$$

Here μ_i^* is the expected value of policymaker i ’s experimental payoff given the first-period outcomes, and 0 is the expected payoff from the sure policy.

¹⁰While a regional policymaker typically learns more from his own experiments, this need not be the case. For example, while $\sigma_1^{*2}[1,0] > \sigma_1^{*2}[0,1]$ typically holds, there are exceptions. One such exceptional case is when $\rho \rightarrow \pm 1$ and $\sigma_2^2 > \sigma_1^2$. Here 2’s experiment is more informative since the policymakers’ payoffs are so tightly connected, and 2’s experiment is more likely to generate extreme results (and so shift 1’s mean term). A formal proof of this case follows from comparing (24) and (25) in Appendix A.1.

The first-period choice is more complicated. The experimental policy provides useful information for the second period, and there is strategic interaction due to the informational externality between regions. Recall from (4) and (5) that given the initial beliefs, μ_i^* is normally distributed with mean μ_i and variance $\sigma_i^{*2}[a_{11}, a_{21}]$. The expected second-period payoff for policymaker i is the probability the experimental policy will be the better choice times the expected experimental payoff,

$$\begin{aligned}
 V_i^{t=2}(a_{11}, a_{21}) &\equiv E\Pr(\mu_i^* > 0 | (a_{11}, a_{21}))E(\mu_i^* | (a_{11}, a_{21}), \mu_i^* > 0) \\
 &= \mu_i \Phi\left(\frac{\mu_i}{\sigma_i^*[a_{11}, a_{21}]}\right) + \sigma_i^*[a_{11}, a_{21}]\phi\left(\frac{\mu_i}{\sigma_i^*[a_{11}, a_{21}]}\right) \quad (8)
 \end{aligned}$$

The expectation operator E is defined with respect to the information available at the start of period one, $\Phi(x)$ is the standard normal cumulative distribution and $\phi(x)$ is the standard normal density.¹¹ Notice that (a_{11}, a_{21}) influences $V_i^{t=2}(\bullet)$ only through $\sigma_i^*[a_{11}, a_{21}]$.

The expected second-period payoff has some important properties and interpretations. It is non-negative ($V_i^{t=2}(\bullet) \geq 0$), increasing in the experimental mean ($\partial V_i^{t=2} / \partial \mu_i > 0$), and increasing in the variance of the second-period experimental mean ($\partial V_i^{t=2} / \partial \sigma_i^*[\bullet] > 0$).¹² The reason a higher mean and variance are advantageous is that (8) is the expected value of a left truncated normal distribution. A large variance implies extreme values are more likely, and negative values are truncated because the sure policy can be used in the second period. This logic suggests interpreting $V_i^{t=2}(\bullet)$ as an expected learning measure. The $\sigma_i^*[\bullet]$ comparative static also helps illustrate the informational externality of experimenting: when $-i$ experiments, $\sigma_i^*[\bullet]$ (and thus $V_i^{t=2}(\bullet)$) increases due to Remark 1.

Conditional on his neighbor's action, policymaker 1's expected first-period payoff is,

$$\begin{aligned}
 V_1^{t=1}(a_{11} = 1, a_{21}) &= \mu_1 + \beta(a_{21} V_1^{t=2}(1, 1) + (1 - a_{21}) V_1^{t=2}(1, 0)) \\
 V_1^{t=1}(a_{11} = 0, a_{21}) &= 0 + \beta(a_{21} V_1^{t=2}(0, 1) + (1 - a_{21}) V_1^{t=2}(0, 0)) \quad (9)
 \end{aligned}$$

The expected payoff for 2 is written in an analogous fashion though reversing the arguments in $V_i^{t=2}(\bullet)$. These formulae partition the expected payoff into a current flow (the first term) and an informational benefit (the second term). One assumption is needed before solving for the optimal first-period action.

¹¹The first equation in (8) uses $E\epsilon_{i2} = 0$ and the 0 payoff of the sure policy. The second equation follows because μ_i^* has a normal distribution (see Johnson and Kotz 1970).

¹²This first point follows from L'Hopital's rule while the remaining two points follow from simple differentiation and the fact that $\phi'(x) = -x\phi(x)$.

ASSUMPTION 1 (FREE-RIDING): *The parameters μ_i , σ_i and ρ satisfy for $i = 1, 2$,*

$$V_i^{t=2}(1, 1) - V_i^{t=2}(0, 1) < V_i^{t=2}(1, 0) - V_i^{t=2}(0, 0) \tag{10}$$

Assumption (1) holds for nearly all parameter values.¹³ It means that policymaker i 's net benefit from experimenting decreases when his neighbor experiments, i.e., $V_i^{t=1}(a_{11} = 1, a_{21}) - V_i^{t=1}(a_{11} = 0, a_{21})$ is decreasing in a_{21} . This condition implies that experiments are strategic substitutes. That is, the policymakers will free-ride off each other's experiments when Assumption 1 is satisfied.

Because each equation in (9) is independent of a_{11} , policymaker 1's optimal first-period action is determined by a simple cut-off rule. Since an analogous condition holds for policymaker 2, the general best response function for policymaker i can be written as,

$$a_{i1} = \begin{cases} 1 & a_{-i1} < a_{-i1}^* \\ a \in [0, 1] & a_{-i1} = a_{-i1}^* \\ 0 & a_{-i1} > a_{-i1}^* \end{cases} \tag{11}$$

where

$$a_{21}^* \equiv \frac{\mu_1 + \beta(V_1^{t=2}(1, 0) - V_1^{t=2}(0, 0))}{\beta(V_1^{t=2}(1, 0) + V_1^{t=2}(0, 1) - V_1^{t=2}(1, 1) - V_1^{t=2}(0, 0))} \tag{12}$$

a_{11}^* is written analogously with $V_2^{t=2}(0, 1)$ replacing $V_1^{t=2}(1, 0)$ in the numerator. (11) states that a policymaker will experiment with probability one if his neighbor is unlikely to experiment, while he uses the sure policy if his neighbor is likely to experiment. In the middle case of (11) the policymaker is indifferent between the two actions and so is willing to mix.

I will focus on the case where $a_{i1}^* \in (0, 1)$ which means that neither policy-maker has a dominant action.¹⁴ This is not only the theoretically most interesting scenario but is also the case with greatest real-world relevance. This is because the experimental policy is unlikely to dominate the sure policy (because then everyone would already be using the experimental policy) and is also unlikely to be dominated by the sure policy (because then no one would be considering the experimental policy).

¹³There are two conditions under which Assumption 1 may fail, and only the first causes any complications. First, the assumption does not hold in a symmetric neighborhood around $\rho = 0$. This neighborhood is quite small: when $\beta = 0.75$, $\mu_1 = -0.1$, $\sigma_i = 1$, then (10) fails if $|\rho| < 0.017$. Second, the assumption does not hold when $\sigma_i \rightarrow 0$ and $\mu_i < 0$. In this case policymaker i will never experiment, there is no strategic interaction, and so the free-rider effect is irrelevant. A formal derivation and further discussion of these points is available upon request.

¹⁴When the experimental policy is dominant for policymaker 1, then under Assumption 1 $a_{21}^* \geq 1$ or $\mu_1 + \beta V_1^{t=2}(1, 1) \geq \beta V_1^{t=2}(0, 1)$. When the sure policy is dominant for policymaker 1, then under Assumption 1 $a_{21}^* \leq 0$ or $\mu_1 + \beta V_1^{t=2}(1, 0) \leq \beta V_1^{t=2}(0, 0)$.

ASSUMPTION 2 (NO DOMINANT ACTION): *The parameters μ_i , σ_i , ρ and β satisfy $a_{i1}^* \in (0,1)$ for $i = 1,2$.*

The assumption holds so long as $\mu_i \in [\mu_i^L(\beta, \sigma_i), \mu_i^H(\beta, |\rho|, \sigma_1, \sigma_2)]$ for some $\mu_i^L \leq \mu_i^H < 0$ with $\partial \mu_i^H / \partial |\rho| > 0$. Assumption 2 is used only to simplify the exposition in the proofs. I discuss below how the key result is modified when it is violated.

In equilibrium (11) must be satisfied for both policymakers. If Assumption 2 holds then there will be a mixed action equilibrium with experimenting probabilities a_{i1}^* .¹⁵ There are several interesting comparative statics of the mixed action equilibrium.

Remark 2 (Comparative Statics): Under Assumptions 1–2,

1. Mixed Action Comparative Statics: $\partial a_{i1}^* / \partial |\rho| < 0$, $\partial a_{i1}^* / \partial \mu_i = 0$, $\partial a_{i1}^* / \partial \mu_{-i} > 0$.
2. Perverse Comparative Statics: An increase in μ_i can strictly decrease the level of aggregate experimentation when it induces a shift from a mixed action to a pure action equilibrium (or vice versa if Assumption 2 does not hold).

Proof: See Appendix A.2.

Under mixed actions an increase in $|\rho|$ decreases each policymaker's propensity to experiment. This is because as the regions become more similar or dissimilar, the policymakers are able to learn more from their neighbor's experiment. Experimentation is therefore reduced because of Assumption 1. Also under mixed actions, an increase in μ_i only influences $-i$'s actions as illustrated in Figure 1.¹⁶ However, an increase in μ_i may make the experimental policy the dominant choice for policymaker i . In this case i uses the experimental policy with probability one and $-i$ free-rides and does not experiment. This reduces aggregate experimentation. Analogous reasoning can be applied when an increase in μ_i initiates a shift from just region $-i$ experimenting to a mixed action equilibrium. These cases are illustrated in Figure 2.

It is important to see whether the decentralized outcome is efficient. The *social optimum* is a strategy for each region which maximizes the total

¹⁵The mixed action equilibrium is not unique but is also compatible with two pure action equilibria where exactly one policymaker experiments (see Figure 1). All of the Propositions still hold if instead the pure strategy equilibria are considered. This is because these results rely on characterizing the conditions for centralized experimentation, and this outcome yields higher experimentation than any of the equilibria above. Assumption 2 is ruling out *unique* equilibria where both, one, or neither policymaker experiments.

¹⁶A policymaker's mixing level is set to make his neighbor indifferent between the two policies. An increase in μ_i makes the experimental policy relatively more attractive to policymaker i . Because of Assumption 1, policymaker $-i$ must increase his probability of experimenting to keep i indifferent.

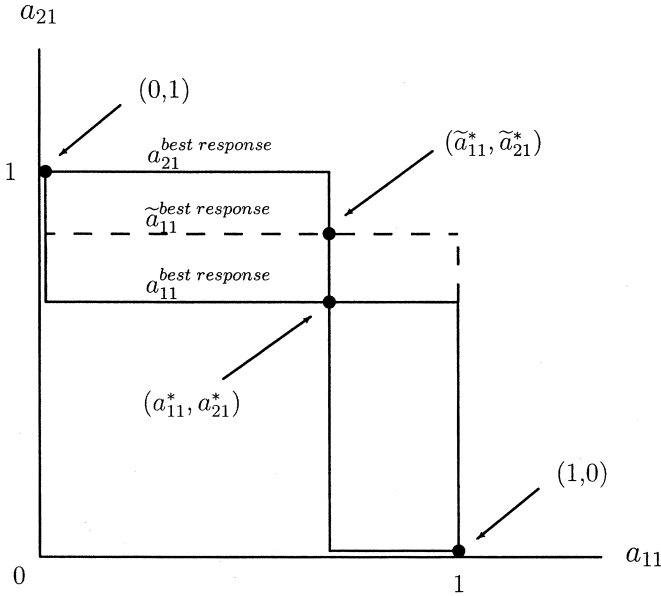


Figure 1: Mixed Action Decentralized Equilibrium

Comments:

The diagram assumes Assumptions 1 and 2 hold. $\tilde{a}_{11}^{bestresponse}$ (- - -) involves a higher μ_1 than $a_{11}^{bestresponse}$.

expected discounted payoff of the region policymakers. The relevant objective function is,

$$\begin{aligned}
 W = \max_{a_{11}, a_{21}} & (a_{11}\mu_1 + a_{21}\mu_2) + \beta[a_{11}a_{21}\mathbf{V}^{t=2}(1,1) + a_{11}(1 - a_{21})\mathbf{V}^{t=2}(1,0) \\
 & + (1 - a_{11})a_{21}\mathbf{V}^{t=2}(0,1) + (1 - a_{11})(1 - a_{21})\mathbf{V}^{t=2}(0,0)] \quad (13)
 \end{aligned}$$

where $\mathbf{V}^{t=2}(a_{11}, a_{21}) \equiv V_1^{t=2}(a_{11}, a_{21}) + V_2^{t=2}(a_{11}, a_{21})$. The social optimum differs from the decentralized outcome in that it explicitly takes into account the learning externality. This is because (13) includes payoffs for both policymakers. As long as pure actions are used, decentralization involves sub-optimal experimentation.¹⁷

Remark 3 (Decentralized Under-Experimentation): Suppose that only pure actions are allowed under decentralization and Assumption 1 holds. Then the

¹⁷This result does not always hold when decentralized mixed actions are used. Under certain parameter values, the social optimum involves one region experimenting while the decentralized outcome involves each region using the experimental policy with probability near one. Decentralization then has a higher level of aggregate experimentation. Figure 3 illustrates this case in the portion where ρ is near ± 0.5 .

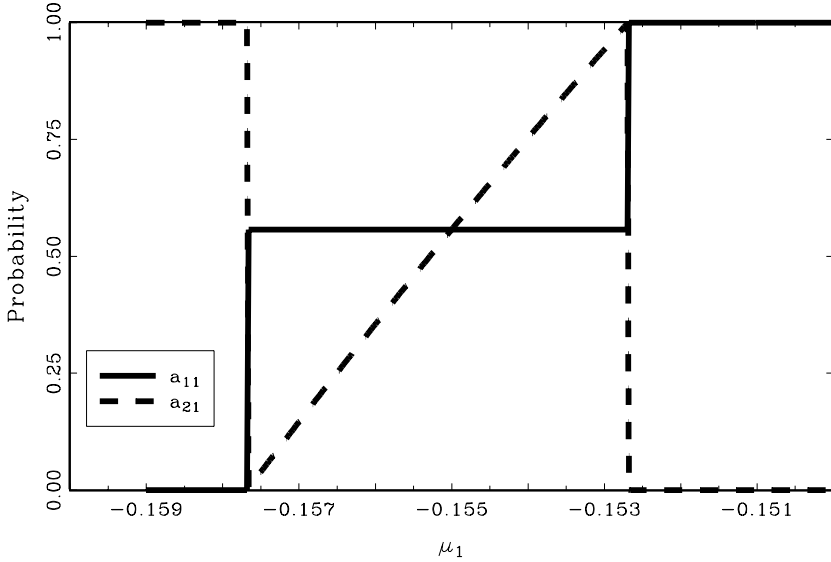


Figure 2: μ_1 Comparative Static ($\beta = 0.75, \rho = 0.2, \mu_2 = -0.155, \sigma_1 = \sigma_2 = 1$)

decentralized outcome weakly exhibits aggregate under-experimentation relative to the social optimum.

Proof: See Appendix A.3.

The intuition is that each decentralized policymaker ignores the informational benefit his experiment provides to his neighbor.¹⁸

For policy purposes it is interesting to consider schemes for decentralizing the social optimum. Presume that a central planner knows all of the relevant parameters but is unable to directly change them. However, he can impose a set of subsidies or taxes on the regions conditioned on their policy choices. Presume also that the subsidies and taxes must exactly balance. If the parameters imply under-experimentation, then a simple tax on first-period actions can implement the social optimum,

$$tax_i(a_{11}, a_{21}) = \begin{cases} \tau & a_{i1} < a_{-i1} \\ -\tau & a_{i1} > a_{-i1} \\ 0 & a_{i1} = a_{-i1} \end{cases} \quad (14)$$

¹⁸In the second period the decentralized and social optimum problems are identical. Any learning due to second-period experiments is irrelevant, since this is the end of the game. Therefore the socially optimal second-period actions are also characterized by (7).

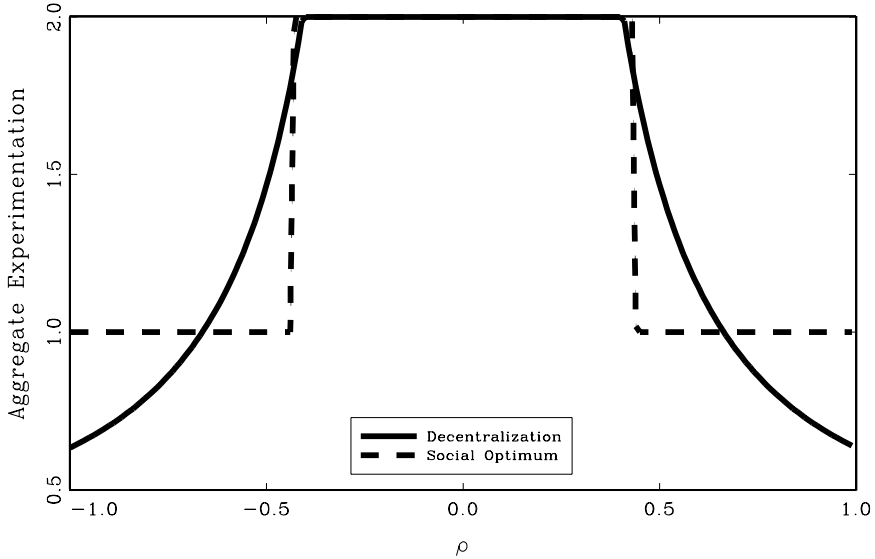


Figure 3: Decentralized Over-Experimentation ($\beta = 0.75$, $\mu_1 = \mu_2 = -0.125$, $\sigma_1 = \sigma_2 = 1$)

where $\tau > 0$.¹⁹ This scheme lowers the sure policy payoff when the other policymaker is experimenting and is like a (conditional) increase in μ_i . By appropriately setting τ it is always possible to induce the social optimum.²⁰ Notice that (14) need not involve actual transfers in equilibrium, but rather it is the threat of a penalty that induces policymakers to pick the appropriate action.

3.3. Centralization

In the *centralized outcome* a single policymaker maximizes the total expected present discounted payoff of the two regions as in (13). However, he is restricted each period to selecting a single policy which is used in both regions. The backwards induction solution is analogous to the decentralized case. The optimal second period action follows a cut-off rule based on the expected return from experimenting,

¹⁹If there is over-experimentation, τ in (14) is replaced with $-\tau$. If there is socially optimal experimentation, then $\tau = 0$ is used.

²⁰When the social optimum has both policymakers experimenting, setting $\tau = \max_i |\mu_i|$ induces socially optimal actions. When the social optimum has one policymaker experimenting, one can show that there exists a range of τ values that induce socially optimal actions. Finally, recall from Appendix A.3 that the social optimum cannot involve mixed actions.

$$a_2 = \begin{cases} 1 & \mu^* \geq 0 \\ 0 & \mu^* < 0 \end{cases} \quad (15)$$

Here $\mu^* \equiv \mu_1^* + \mu_2^*$ is the expected experimental return in the second period. For a given first-period action, the optimal second-period payoff is,

$$\begin{aligned} \tilde{V}^{t=2}(a_1) &\equiv E\Pr(\mu^* > 0 | a_1)E(\mu^* | a_1, \mu^* > 0) \\ &= \mu\Phi\left(\frac{\mu}{\sigma^*[a_1]}\right) + \sigma^*[a_1]\phi\left(\frac{\mu}{\sigma^*[a_1]}\right) \end{aligned} \quad (16)$$

The second equality follows because μ^* has a normal distribution with mean $\mu \equiv \mu_1 + \mu_2$ and variance,

$$\sigma^{*2}[a_1] \equiv a_1(\sigma_1^{*2}[1,1] + \sigma_2^{*2}[1,1] + 2\text{Cov}[1,1]) \quad (17)$$

The $\text{Cov}[1,1]$ term is the covariance between the μ_i^* 's when both regions experiment (this is the off-diagonal term in the covariance matrix from (5) when α is the identity matrix). The optimal first period action is determined by,

$$\tilde{V}^{t=1} = \max_{a_1} a_1\mu + \beta\tilde{V}^{t=2}(a_1) \quad (18)$$

This means that the optimal first-period action is to experiment, $a_1 = 1$, if and only if $\mu + \beta\tilde{V}^{t=2}(1)$ is non-negative. Notice that only pure actions could be optimal.

The chief weakness of centralization is that it is restricted to uniform policies. When the regions are quite different, this will be a serious limitation because differential policies are more appropriate. However, the main strength of centralization is that it internalizes all of the learning externalities of experimenting. This advantage is particularly important when the regions are similar so the learning externality is large.

As this reasoning suggests, the comparison with the social optimum is a bit complicated.

Remark 4 (Centralized Over- and Under-Experimentation): The centralized outcome may exhibit aggregate over- or under experimentation relative to the social optimum.

While the centralized outcome and the social optimum share the same ex ante objective function (13), only the social optimum can assign different actions to the two regions. When the regions are quite different, the social optimum may involve one region experimenting while the centralized outcome must have neither or both regions experimenting. That is, a very high or low expected experimental payoff in one region may

swamp the expected return in the other region.²¹ Because of this complication, I will not focus on the social optimum in the remainder of the text.

3.4. The Main Result

The main objective of this article is to compare aggregate experimentation under centralization and decentralization.

PROPOSITION 1: *Under Assumptions 1–2, there exists a $\rho^* > 0$ such that aggregate experimentation is higher under centralization than decentralization if and only if $\rho > \rho^*$.*

Proof: See Appendix A.4.

To understand this result, recall that Assumption 2 implies that decentralization involves mixed actions in the first period. Because centralization involves pure actions, it will be sufficient to consider the conditions that imply centralized experimentation. A first-period centralized experiment has both a current payoff and a learning benefit. When ρ is positive, higher ρ values mean the regions are becoming more similar. This increases both the current payoff (since uniform policies are more appropriate) and the learning benefit (since it becomes easier to disentangle the permanent and idiosyncratic payoff components). Alternatively, when ρ is negative then higher values mean the regions are becoming less dissimilar (or, equivalently, closer to being independent). This increases the current payoff but *decreases* the learning benefit (since learning would be higher if the regions were becoming *more* dissimilar). The first effect dominates which means that higher values of ρ always increase the centralized payoff from experimenting. Hence the centralized outcome involves experimenting only when ρ is large enough.²²

This result and the others in the remainder of the article can be extended to allow for dominant actions (Assumption 2 fails).²³ When the experimental policy is dominant for one or more regions, Proposition 1 still holds though ρ^* is smaller and possibly negative. When the sure policy is dominant for one or more regions, Proposition 1 still holds as

²¹One example is $\mu_1 \rightarrow -\infty$ and $\mu_2 = \underline{\mu}_2$ where $\underline{\mu}_2$ is defined in Appendix A.4. In this case the social optimum has region 2 experimenting while the centralized outcome has no experimenting at all.

²²More formally, ρ influences the centralized payoff through its effect on $\sigma^{*2}[1]$ defined in (17). There are two channels of influence: higher ρ values increase $\text{Cov}[1,1]$, and higher *absolute* ρ values increase $\sigma_i^{*2}[1,1]$. These channels imply that $\sigma^{*2}[1]$ is increasing in ρ when $\rho \geq 0$. Appendix A.4 shows that $\sigma^{*2}[1]$ is *always* increasing in ρ . Because $\bar{V}^{t=2}(\cdot)$ (and thus $\bar{V}^{t=1}$) increases in $\sigma^{*2}[1]$, centralized experimentation will only occur if ρ is relatively large.

²³It is harder to extend the result when Assumption 1 fails, since then decentralized experiments are strategic complements. In this case the pure strategy equilibria involve neither *or* both regions experimenting (under Assumption 2), and it is difficult to argue which outcome is more reasonable.

long as neither or both μ_i 's are extremely negative.²⁴ The one case where the Proposition cannot be extended is when only one μ_i is extremely negative. Here there is no centralized experimentation (because $\mu_i \ll 0$ predominates) while policymaker $-i$ typically experiments under decentralization. Experimentation is unambiguously higher under decentralization.

Proposition 1 focuses on how ρ influences centralized experimentation. The result can be extended to consider the *difference* in aggregate experimentation between centralization and decentralization.²⁵

COROLLARY 1: *Under Assumptions 1–2, when $\rho \geq 0$ the difference in aggregate experimentation between centralization and decentralization is increasing in ρ .*

Appendix A.4 shows that the centralized benefit of experimenting is always increasing in ρ . This means that centralized experimentation is weakly increasing in ρ . Corollary 1 follows because higher absolute values of ρ decrease decentralized experimentation through the free-rider effect (see Remark 2). This result again suggests that greater regional similarity tends to promote centralized experimentation relative to decentralized experimentation.

4. Extensions

4.1. Arbitrary Number of Regions

In this section I consider how the comparison between decentralization and centralization is altered by allowing for an arbitrary number of regions, multiple experimental policies, and political yardstick competition.²⁶ To begin to say that there are now $N \geq 2$ regions indexed by i . For analytical tractability suppose that the regions have a common permanent experimental variance, $\sigma_i^2 = \sigma^2$, and a common experimental correlation, ρ . This means the $N \times N$ variance matrix in (3) can be written as,

$$\Sigma_N \equiv \begin{pmatrix} 1 & \rho & \rho & \dots \\ \rho & 1 & \rho & \dots \\ & & \ddots & \\ \rho & \rho & \dots & 1 \end{pmatrix} \sigma^2 \quad (19)$$

²⁴The first result follows since the proof in Appendix A.4 uses Assumption 2 only to show $\rho^* > 0$. The second result follows if both μ_i 's are not much lower than the $\underline{\mu}_i$'s defined in Appendix A.4 (this only slightly increases $\rho^* < 1$) or if both μ_i 's are quite negative (then decentralization involves no experimentation).

²⁵The restriction $\rho \geq 0$ in Corollary 1 is quite reasonable in practice. It rules out the case that an experimental success in one region signals a likely experimental failure in another region.

²⁶The model can also be extended to allow the idiosyncratic component of the experimental payoffs to be correlated across regions. This correlation could reflect temporary phenomena that influence all regions, such as business cycles or taste fads. Details of this extension are contained in an earlier version of the paper which is available upon request.

To ensure Σ_N is well defined, $\rho \geq -(N-1)^{-1}$ must be imposed.²⁷ Because of the common variance and correlation assumption, each neighbor's experiment has a symmetric effect. That is, from the first-period perspective a regional policymaker only cares about the total number of his neighbors who use the experimental policy. However, there is still ex ante heterogeneity (since the μ_i 's can be different) and ex post heterogeneity (since the regions can have different first-period outcomes).

Assumption 1 needs to be extended. The new condition states that a policymaker's net benefit of experimenting is decreasing in the number of his neighbors who experiment. More formally, let policymaker i 's expected second period payoff when M neighbors experiment be $V_i^{t=2}(a_{i1}, M)$.

ASSUMPTION 3 (FREE-RIDING WITH N REGIONS): *The parameters μ_i , σ , and ρ are such that $V_i^{t=2}(1, M) - V_i^{t=2}(0, M)$ is decreasing in M .*

As with Assumption 1, numerical solutions indicate that the condition holds for almost all parameter values. I will continue to assume that no policymaker has a dominant action (Assumption 2) so decentralization involves mixed actions.²⁸

PROPOSITION 2: *Suppose there are $N \geq 2$ regions. Under Assumptions 2–3 there exists a $\rho^*(N) > 0$ with $\partial \rho^*(N)/\partial N < 0$ such that aggregate experimentation is higher under centralization than decentralization if and only if $\rho > \rho^*(N)$.*

Proof: See Appendix A.5.

The intuition for Proposition 2 is similar to the two-region case. A higher ρ encourages centralized experimentation because it increases the potential reward from a successful first-period experiment. When N increases a centralized experiment in the first period is more informative, since it becomes easier to disentangle the (independent) idiosyncratic and (correlated) permanent components. This means that the parameter space under which centralization involves more experimentation than decentralization grows with N .²⁹

Proposition 2 can also be extended to consider the difference in aggregate experimentation between centralization and decentralization. As N grows, the decentralized free-riding effect becomes stronger. This is because each regional policymaker only cares about the number of other policymakers who experiment. When more neighbors experiment, the

²⁷A variance matrix must be positive semidefinite which requires that all of the principal minor determinants are non-negative. Some algebra shows this is equivalent to $\Delta_N \geq 0$ where Δ_N is defined below (47). Some further algebra yields the condition in the text.

²⁸When the experimental policy is dominant for policymaker i , then under Assumption 3 he is willing to experiment even if all of his neighbors do, $\mu_i + \beta V_i^{t=2}(1, N-1) \geq \beta V_i^{t=2}(0, N-1)$. If the sure policy is dominant for policymaker i , then under Assumption 3 he is unwilling to experiment even if none of his neighbors do, $\mu_i + \beta V_i^{t=2}(1, 0) \leq \beta V_i^{t=2}(0, 0)$.

²⁹The restriction $\rho \geq -(N-1)^{-1}$ rules out the possibility of more centralized experimenting when ρ becomes quite negative (e.g., the learning benefit grows as ρ approaches -1).

incentive to experiment decreases due to Assumption 3. Because of Proposition 2, this means that the difference between centralized and decentralized experimentation grows with N .

COROLLARY 2: *Suppose there are $N \geq 2$ regions. Under Assumptions 2–3, the difference in aggregate experimentation between centralization and decentralization is increasing in N and when $\rho \geq 0$ is increasing in ρ .*

4.2 Multiple Experimental Policies

The next extension adds a second experimental policy option. This generalization is important because policies have multiple characteristics, and it is possible to experiment along each characteristic.³⁰ For analytical simplicity again assume there are two regions. The two experimental policies, labeled $j = A, B$, are identical in the sense that they have the same initial experimental mean, variance, and correlation for the two regions. The common mean terms are $\mu_{ij} = \mu$, the variance terms are $\sigma_{ij}^2 \equiv \sigma^2$, and the correlation terms are $\rho_j = \rho$ for $i = 1, 2, j = A, B$.³¹ Let a_{ijt} be the action of region i on policy j in period t . Denote the expected second-period payoff in region i as $V_i^{t=2}(a_{1A1}, a_{2A1}; a_{1B1}, a_{2B1})$. In the second-period decentralized policymaker i uses whichever policy gives the highest expected return,

$$V_i^{t=2}(a_{1A1}, a_{2A1}; a_{1B1}, a_{2B1}) = \max_{a_{ij2}} E(\mu_{iA}^*, \mu_{iB}^*; 0) \tag{20}$$

The values on the right-hand side are the expected payoffs from the two experimental policies (conditional on the first-period actions) and the sure policy. Under centralization, the i subscript is omitted and the experimental return is summed over the regions, $\mu_j^* \equiv \mu_{1j}^* + \mu_{2j}^*$.

The addition of the second experimental policy does not change the centralized choice so long as $\mu_j \equiv \mu_{1j} + \mu_{2j}$ is negative. This is because the centralized policymaker is restricted to one policy per period, and he is indifferent between the ex ante identical experimental policies. If he does not use experimental policy j in the first period, he does not learn about it and will not use it in the second period when $\mu_j < 0$. For the two decentralized policymakers, however, the experimental policies might be strategic complements. Each policymaker may use a different experimental policy because this allows him to learn a bit about both policies. It is important to stress that each policymaker is acting in a self-interested fashion and still does not take into account the potential learning externality. Rather, an experiment becomes less appealing to a policymaker when his neighbor is using it. This is because part of the benefit of

³⁰For example, welfare policies can have a maximum tenure length as well as a work requirement.

³¹The results in this section can be extended to more general parameter values.

experimenting with policy j is learning about its distribution, and a neighbor's experiment reduces the potential for learning through the free-rider effect. Instead, it may be more informative to experiment with the other policy.

The definition of aggregate experimentation must be modified to account for the additional experimental policy. The measure should still capture the expected learning from first-period actions. Recall that $V_i^{t=2}(\bullet)$ is a learning measure because it reflects how much beliefs are expected to shift due to the first-period actions. Therefore define *aggregate experimentation with multiple policies* to be $V_1^{t=2}(\bullet) + V_2^{t=2}(\bullet)$.³² Notice that this reduces to the original definition when there is only one experimental policy.

PROPOSITION 3: *Suppose there are two ex ante identical experimental policies. Aggregate experimentation may be higher under decentralization than centralization even if $\rho \rightarrow 1$.*

Proof: See Appendix A.6.

Even when centralization involves no experimentation, the decentralized policymakers may each use different experimental policies. This is shown formally in Appendix A.6 for the case $\rho = -1$. The intuition is that negative correlation induces a decentralized free-rider effect on a *particular* experimental policy because a policymaker's experiment is quite informative for his neighbor. Alternatively, a negative correlation means an experimental policy is unlikely to be effective in both regions which reduces the centralized benefit of experimenting.

This extension provides a somewhat surprising explanation for Justice Brandeis's conjecture that decentralization involves policy diversity. Each regional policymaker may choose a unique experiment, since this allows him to learn more than he would from simply matching his neighbor's experiment. That is, *free-riding* discourages use of the same policy and (when multiple experiments are available) promotes decentralized policy diversity.

4.3 Political Yardstick Competition

It is sometimes argued that decentralization involves greater policy innovation because of political competition between regions.³³ The idea is that regional policymakers are interested in running for higher office, and one way to signal efficacy to voters is to be the first to devise an innovative

³²An equivalent definition is to sum each region's maximum second-period expected experimental variance.

³³For example, in a discussion on welfare policy Wisconsin governor Tommy Thompson argued, "But right now you've got governors who are so darn competitive. They don't want to read that they're not taking care of the poor, and they're not going to let a governor in an adjoining state get ahead of them. We're very competitive" (*New York Times*, 5 July 1998).

policy. Besley and Case (1995) present one such model in which voters evaluate their local policymaker based on the relative success of his fiscal policies compared to those of his neighbors. They refer to this as “yardstick competition” because each policymaker is evaluated using his neighbors’ performance as the benchmark. This kind of race to introduce policy experiments might be absent at the central level because the chief beneficiary of a centralized innovation, the central policymaker, is unlikely to have further political aspirations.

While such yardstick competition undoubtedly exists, whether it promotes decentralized policy innovation depends on the voter evaluation process. Assume that the regional payoffs discussed in previous sections are some measure of regional voter welfare. Presume that the regional policymakers are chiefly concerned with maximizing their region’s welfare, but they also have personal political ambitions. A policymaker receives a payoff bonus, $G > 0$, if he wins the yardstick competition with his neighbor. Assume first that the winner is the policymaker with the most successful experimental outcome in the first period. Policymaker 1’s payoff from (9) becomes,

$$\begin{aligned} \bar{V}_1^{t=1}(a_{11} = 1, a_{21}) &= V_1^{t=1}(a_{11} = 1, a_{21}) \\ &+ \beta(a_{21} \Pr(\pi_{11} > \pi_{12})G + (1 - a_{21})G) \\ \bar{V}_1^{t=1}(a_{11} = 0, a_{21}) &= V_1^{t=1}(a_{11} = 0, a_{21}) \end{aligned} \quad (21)$$

The payoff from experimenting (the top line) has increased because of the possibility of winning the yardstick competition. The first G term represents the probability of 1 winning when both experiment while the second G term is 1’s guaranteed bonus when he is the only experimenter. The payoff from not experimenting (the bottom line) is unchanged. Because similar incentives hold for policymaker 2, it is not difficult to see that this sort of yardstick competition induces weakly higher aggregate experimentation.³⁴

Now suppose that the yardstick winner is the policymaker with the highest first-period outcome regardless of which action is used. Such a reward scheme would result if voters are myopic and only care about current policy outcomes (and ignore the future benefits of learning). This means that a policymaker can receive a bonus when he uses the sure policy if his neighbor’s experiment performs poorly or if his neighbor also selects the sure policy. Policymaker 1’s payoff from (9) is now,

³⁴A similar result holds (i) if the bonus is only awarded when the experimental outcome is positive; or (ii) if the bonus is awarded to the policymaker whose experimental outcome relative to expectations, $\pi_{i1} - \mu_i$, is largest.

$$\begin{aligned} \bar{V}_1^{t=1}(a_{11} = 1, a_{21}) &= V_1^{t=1}(a_{11} = 1, a_{21}) \\ &\quad + \beta \left(a_{21} \Pr(\pi_{11} > \pi_{12}) G + (1 - a_{21}) \Phi \left(\frac{\mu_1}{\sigma_1^* [1, 0]} \right) G \right) \\ \bar{V}_1^{t=1}(a_{11} = 0, a_{21}) &= V_1^{t=1}(a_{11} = 0, a_{21}) \\ &\quad + \beta \left(a_{21} \Phi \left(\frac{-\mu_2}{\sigma_2^* [0, 1]} \right) G + (1 - a_{21}) 0.5 G \right) \end{aligned} \quad (22)$$

When policymaker 1 experiments (the top line), he wins the competition if his experiment is more successful than his neighbor's experiment (the first G term) or than his neighbor's sure policy (the second G term). When he uses the sure policy (the bottom line), he wins if his neighbor's experiment is unsuccessful (the first G term), and he splits the bonus if neither experiment (the second G term). When experimenting is non-dominant, $\mu_i < 0$ and mixed actions are used.³⁵ Some algebra shows that both policymakers place a smaller weight on experimenting relative to the case without yardstick competition. The intuition is that the immediate payoff from the sure policy is higher than from the experimental policy, and so a policymaker is more likely to win G if he uses the sure policy. It is therefore impossible to determine whether yardstick competition increases decentralized policy experimentation without specifying the voter evaluation process.

5. Conclusion

In this paper I investigate the role of government decentralization in the policy innovation process. Contrary to conventional wisdom, local governments need not be better innovators than the central government. While each local government can simultaneously consider a different policy, they each ignore the external benefit of their choice: when any one government experiments with a policy, all governments learn about its potential feasibility. Local governments may forego experimenting and free-ride off the experience of their neighbors. In contrast the central government is able to internalize these informational externalities but is restricted to using a single policy at any moment. Centralization results in greater experimentation when the local governments are relatively homogeneous or large in number, but decentralization may have greater experimentation when there are multiple policy options available.

³⁵When experimenting is dominant, then the yardstick competition described here has no effect.

This paper can help analyze how the recent U.S. welfare reform (see Blank 1997 for details) may influence policy innovation. The shift of federal aid from matching to block grants might discourage experimentation. This is because matching but not block grants increase with the level of state spending, and experimental policies such as Wisconsin's training program tend to be quite expensive (see Wiseman 1996). Alternatively, the welfare reform also streamlined or eliminated the time-consuming waiver applications which states formerly had to complete in order to use a policy experiment. This change is likely to induce more innovation. It would be interesting to formally evaluate such changes in a model that adds institutional detail to the framework presented here.

This analysis also has empirical implications. Previous empirical work on decentralized policy innovation has ignored the role of inter-state or inter-region heterogeneity. A reasonable strategy for future research would be to measure the correlation across states in some objective policy outcome. For example, one could calculate the correlation across states in poverty or single motherhood rates when welfare term limits are imposed. Another example is to see whether community policing has a uniform effect on crime rates. These data could then be used to test implications of the model, such as the prediction that greater inter-state heterogeneity increases the rate of innovation under decentralization.

This paper makes the simplifying assumption that centralization and decentralization differ only in whether learning externalities are internalized and whether differentiated policies can be implemented. The model can be extended to consider other potential differences. First, the assumption that decentralized policy payoffs are independent across regions could be relaxed. In reality, local policy outcomes are likely to be inter-related, say, because one state's generous welfare program attracts migrants from less generous states. Decentralized experimentation is likely to be relatively larger (smaller) if such externalities are negative (positive). Second, there could be information asymmetries. The central government may have limited information about the preferences of each region though it might be better informed about the degree of similarity between regions. It is not clear whether such informational differences will induce more or less centralized experimentation. Third, the central government may have a more skilled bureaucracy and may enjoy returns to scale in policy research and program evaluation. This suggests the central government might initially select more promising experimental policies which should increase the relative level of centralized experimentation. This also suggests that the central government could learn more from previous experience than do local governments. While this modification should encourage centralized experimentation, it is also likely to increase decentralized experimentation through a reduction in the free-rider effect. Fourth, there may be differences in risk

aversion. The central government may be more inclined to experiment because a bad outcome in one region may be offset by a favorable outcome in a second region. Such implicit insurance is unavailable to local governments. If policymakers have a strong aversion to negative outcomes, this is likely to result in relatively more centralized experimentation. Fully investigating these extensions are interesting topics for future research.

Appendix

A.1 Proof of Remark 1: I will use the formula in (5) to write expressions for $\sigma_1^{*2}[a_{11}, a_{21}]$ (the formulae for $\sigma_2^{*2}[a_{11}, a_{21}]$ are analogous). When neither policy-maker experiments ($a_{11} = 0, a_{21} = 0$), then α is the zero matrix. Substituting this into (5) and solving for the upper left term gives,

$$\sigma_1^{*2}[0,0] = 0 \quad (23)$$

When only policymaker 1 experiments ($a_{11} = 1, a_{21} = 0$), then $\alpha = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$ and so,

$$\sigma_1^{*2}[1,0] = \frac{\sigma_1^4}{\sigma_1^2 + 1} \quad (24)$$

When only policymaker 2 experiments ($a_{11} = 0, a_{21} = 1$), then $\alpha = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$ and so,

$$\sigma_1^{*2}[0,1] = \frac{\rho^2 \sigma_1^2 \sigma_2^2}{\sigma_2^2 + 1} \quad (25)$$

If both experiment ($a_{11} = 1, a_{21} = 1$), then α is the identity matrix and,

$$\sigma_1^{*2}[1,1] = \frac{((1 - \rho^2)\sigma_1^2 \sigma_2^2 + \sigma_1^2 + \rho^2 \sigma_2^2)\sigma_1^2}{(1 - \rho^2)\sigma_1^2 \sigma_2^2 + \sigma_1^2 + \sigma_2^2 + 1} \quad (26)$$

It is clear that $\sigma_1^{*2}[a_{11}, a_{21}]$ is weakly minimized when neither policymaker experiments. Now,

$$\sigma_1^{*2}[1,1] - \sigma_1^{*2}[1,0] = \frac{\rho^2 \sigma_2^2 \sigma_1^2}{((1 - \rho^2)\sigma_1^2 \sigma_2^2 + \sigma_1^2 + \sigma_2^2 + 1)(\sigma_1^2 + 1)} \quad (27)$$

which must be non-negative and is positive when $\rho \neq 0$. Also,

$$\sigma_1^{*2}[1,1] - \sigma_1^{*2}[0,1] = \frac{[((1 - \rho^2)\sigma_2^2 + 2)(1 - \rho^2)\sigma_2^2 + 1]\sigma_1^4}{((1 - \rho^2)\sigma_1^2\sigma_2^2 + \sigma_1^2 + \sigma_2^2 + 1)(\sigma_2^2 + 1)} \quad (28)$$

which is positive. This shows that $\sigma_1^{*2}[a_{11}, a_{21}]$ is maximized when both policymakers experiment. ■

A.2 *Proof of Remark 2:* I will derive comparative statics for a_{21}^* (the comparative statics for a_{11}^* are analogous). Before beginning, notice that a necessary condition for mixing is that $\mu_i < 0$ which implies $V_i^{t=2}(0,0) = 0$. Differentiating (12) with respect to $|\rho|$ yields,

$$\frac{\partial a_{21}^*}{\partial |\rho|} = \frac{\left[\phi\left(\frac{\mu_1}{\sigma_1^*[1,1]}\right) \frac{\partial \sigma_1^*[1,1]}{\partial |\rho|} - \phi\left(\frac{\mu_1}{\sigma_1^*[0,1]}\right) \frac{\partial \sigma_1^*[0,1]}{\partial |\rho|} \right] a_{21}^* + \phi\left(\frac{\mu_1}{\sigma_1^*[1,0]}\right) \frac{\partial \sigma_1^*[1,0]}{\partial |\rho|} (1 - a_{21}^*)}{V_1^{t=2}(0,1) + V_1^{t=2}(1,0) - V_1^{t=2}(1,1)} \quad (29)$$

using $\partial V_i^{t=2}(a_{11}, a_{21})/\partial \sigma_i^*[a_{11}, a_{21}] = \phi(\mu_i/\sigma_i^*[a_{11}, a_{21}])$. Because the denominator of (29) is positive by Assumption 1, the derivative will have the same sign as the numerator. Differentiation of (24) shows that $\partial \sigma_1^*[1,0]/\partial |\rho| = 0$. The remaining terms in the numerator must be negative or else Assumption 1 will fail. This shows that $\partial a_{21}^*/\partial |\rho| < 0$. Differentiation of (12) shows that $\partial a_{21}^*/\partial \mu_2 = 0$ because $\partial V_1^{t=2}(a_{11}, a_{21})/\partial \mu_2 = 0$. Differentiating (12) with respect to μ_1 yields,

$$\frac{\partial a_{21}^*}{\partial \mu_1} = \frac{1 + \beta \left[\left(\Phi\left(\frac{\mu_1}{\sigma_1^*[1,1]}\right) - \Phi\left(\frac{\mu_1}{\sigma_1^*[0,1]}\right) \right) a_{21}^* + \Phi\left(\frac{\mu_1}{\sigma_1^*[1,0]}\right) (1 - a_{21}^*) \right]}{\beta [V_1^{t=2}(0,1) + V_1^{t=2}(1,0) - V_1^{t=2}(1,1)]} \quad (30)$$

using $\partial V_i^{t=2}(a_{11}, a_{21})/\partial \mu_i = \Phi(\mu_i/\sigma_i^*[a_{11}, a_{21}])$. The numerator is positive (because $\sigma_1^*[1,1] \geq \sigma_1^*[0,1]$) and the denominator is positive (due to Assumption 1), so $\partial a_{21}^*/\partial \mu_1 > 0$.

To show the perverse comparative static, I will consider the case where there is initially a mixed action equilibrium. Presume that μ_1 is at the highest value which is compatible with mixed actions. Policymaker 2 must experiment with probability approaching one to keep policymaker 1 indifferent (as 1 is favorably inclined towards experimenting). Because policymaker 1 is himself mixing, this means the level of aggregate experimentation exceeds one. When μ_1 is increased a bit, policymaker 1 will strictly prefer experimenting. This induces 2 to free-ride and use the sure policy (because of Assumption 1). Hence, an increase in μ_1 lowers the level of aggregate experimentation. A

similar argument can be used to show that an increase in μ_1 will decrease aggregate experimentation when the equilibrium shifts from just policymaker 2 experimenting to mixed actions. ■

A.3 Proof of Remark 3: The first step is to show that only pure actions are used in the first period of the social optimum. I will assume $\mu_i < 0$ because otherwise it is clear that the social optimum involves both policymakers experimenting. Differentiating (13) yields the first order conditions,

$$\begin{aligned} \frac{\partial W}{\partial a_{21}}(a_{21}) &= \mu_1 + \beta[a_{21}\mathbf{V}^{t=2}(1,1) + (1 - a_{21})\mathbf{V}^{t=2}(1,0) - a_{21}\mathbf{V}^{t=2}(0,1)] \\ \frac{\partial W}{\partial a_{21}}(a_{11}) &= \mu_2 + \beta[a_{11}\mathbf{V}^{t=2}(1,1) - a_{11}\mathbf{V}^{t=2}(1,0) + (1 - a_{11})\mathbf{V}^{t=2}(0,1)] \end{aligned} \quad (31)$$

using $\mathbf{V}^{t=2}(0,0) = 0$ when $\mu_i < 0$. (31) can be used to partition the parameter space. The cross-partial derivatives of W are negative due to Assumption 1. This means the conditions for dominant pure actions are $\partial W/\partial a_{i1}(a_{-i1} = 1) \geq 0 \rightarrow a_{i1} = 1$ and $\partial W/\partial a_{i1}(a_{-i1} = 0) < 0 \rightarrow a_{i1} = 0$, and so the following cases involve unique, pure actions,

$$\begin{aligned} \frac{\partial W}{\partial a_{i1}}(a_{-i1} = 1) \geq 0, \quad \frac{\partial W}{\partial a_{-i1}}(a_{i1} = 1) \geq 0 &\rightarrow (a_{i1} = 1, a_{-i1} = 1) \\ \frac{\partial W}{\partial a_{i1}}(a_{-i1} = 1) \geq 0, \quad \frac{\partial W}{\partial a_{-i1}}(a_{i1} = 1) < 0 &\rightarrow (a_{i1} = 1, a_{-i1} = 0) \\ \frac{\partial W}{\partial a_{i1}}(a_{-i1} = 0) < 0, \quad \frac{\partial W}{\partial a_{-i1}}(a_{i1} = 0) \geq 0 &\rightarrow (a_{i1} = 0, a_{-i1} = 1) \\ \frac{\partial W}{\partial a_{i1}}(a_{-i1} = 0) < 0, \quad \frac{\partial W}{\partial a_{-i1}}(a_{i1} = 0) < 0 &\rightarrow (a_{i1} = 0, a_{-i1} = 0) \end{aligned} \quad (32)$$

The only remaining case is when $\partial W(a_{-i1} = 1)/\partial a_{i1} < 0 \leq \partial W(a_{-i1} = 0)/\partial a_{i1}$ for both i . Three action pairs are possible optima (in the sense of meeting the first order Kuhn-Tucker conditions): $(a_{11} = 1, a_{21} = 0)$, $(a_{11} = 0, a_{21} = 1)$, and $(a_{11} = \tilde{a}_{11}, a_{21} = \tilde{a}_{21})$ where $\tilde{a}_{i1} \in (0,1)$ satisfies $\partial W(\tilde{a}_{i1})/\partial a_{-i1} = 0$. The mixed action pair cannot be optimal, since it is an inflection point for the welfare function W (that is, $\partial^2 W/\partial a_i^2 = 0$). Only one of the remaining two pairs—which involve pure actions—can maximize W with ties possible only for parameter values of zero measure. The formal condition is,

$$\frac{\partial W}{\partial a_{i1}} (a_{-i1} = 1) < 0 \leq \frac{\partial W}{\partial a_{i1}} (a_{-i1} = 0)$$

$$\rightarrow \begin{cases} (a_{11} = 1, a_{21} = 0) \\ \mu_1 + \beta \mathbf{V}^{t=2}(1,0) > \mu_2 + \beta \mathbf{V}^{t=2}(0,1) \\ (a_{11} = 1, a_{21} = 0) \text{ or} \\ (a_{11} = 0, a_{21} = 1) \\ \mu_1 + \beta \mathbf{V}^{t=2}(1,0) = \mu_2 + \beta \mathbf{V}^{t=2}(0,1) \\ (a_{11} = 0, a_{21} = 1) \\ \mu_1 + \beta \mathbf{V}^{t=2}(1,0) < \mu_2 + \beta \mathbf{V}^{t=2}(0,1) \end{cases} \quad (33)$$

These results show that the social optimum always involves pure actions.

The next step is to show that decentralization exhibits weak under-experimentation when pure actions must be used. The proof proceeds by using (32) and (33) to partition the parameter space. I will assume that $\mu_i < 0$ because otherwise it is clear that the decentralized and social optima involve both policymakers experimenting with probability one.³⁶ Recall that this inequality implies that $V_i^{t=2}(0,0) = 0$. First presume there is no experimentation in the social optimum ($a_{11} = 0, a_{21} = 0$) which requires that $\partial W(a_{-i1} = 0)/\partial a_{i1} < 0$. Assume the decentralized outcome has both policymakers experimenting which requires that $V_i^{t=1}(a_{i1} = 1, a_{-i1} = 1) \geq V_i^{t=1}(a_{i1} = 0, a_{-i1} = 1) \geq 0$. But because $\partial W(a_{-i1} = 1)/\partial a_{i1} \leq \partial W(a_{-i1} = 0)/\partial a_{i1} < 0$,

$$\begin{aligned} V_1^{t=1}(1,1) + V_2^{t=1}(1,1) &\leq (\mu_1 + \mu_2) + \beta(\mathbf{V}^{t=2}(1,0) + \mathbf{V}^{t=2}(0,1)) \\ &= \frac{\partial W}{\partial a_{11}} (a_{21} = 1) + \frac{\partial W}{\partial a_{21}} (a_{11} = 1) \\ &< 0 \end{aligned} \quad (34)$$

using (9) and (31). The inequality in the top row follows from Assumption 1. Instead assume that one policymaker (say, 1) experiments in the decentralized outcome which requires that $V_1^{t=1}(1,0) \geq V_1^{t=1}(0,0) = 0$. But from (9) and (31),

$$\begin{aligned} V_1^{t=1}(1,0) &= \frac{\partial W}{\partial a_{11}} (a_{21} = 0) - \beta V_2^{t=2}(1,0) \\ &< 0 \end{aligned} \quad (35)$$

³⁶When $\mu_i \geq 0 > \mu_{-i}$, the social optimum has policymaker i experimenting and possibly policymaker $-i$ experimenting if μ_{-i} is not too negative. It is not difficult to show the decentralized optimum exhibits weak under-experimentation in this case.

where the inequality follows since $V_2^{t=2}(1,0) \geq 0$. The case of just policymaker 2 experimenting gives an analogous contradiction. This shows by a process of elimination that the decentralized outcome must have no experimentation when the social optimum has none.

The next case presumes that only policymaker 1 experiments in the social optimum, $(a_{11} = 1, a_{21} = 0)$. The proof for the case $(a_{11} = 0, a_{21} = 1)$ is analogous and is omitted. (32) and (33) show that this case occurs under any of three conditions: (i) $\partial W(a_{21} = 1)/\partial a_{11} \geq 0$, $\partial W(a_{11} = 1)/\partial a_{21} < 0$, (ii) $\partial W(a_{21} = 0)/\partial a_{11} \geq 0$, $\partial W(a_{11} = 0)/\partial a_{21} < 0$, (iii) $\partial W(a_{i1} = 1)/\partial a_{i1} < 0 \leq \partial W(a_{i1} = 0)/\partial a_{i1}$ and $\mu_1 + \beta \mathbf{V}^{t=2}(1,0) \geq \mu_2 + \beta \mathbf{V}^{t=2}(0,1)$. Assume the decentralized outcome has both policymakers experimenting which requires that $V_i^{t=1}(a_{i1} = 1, a_{-i1} = 1) \geq V_i^{t=1}(a_{i1} = 0, a_{-i1} = 1)$. Under social optimum condition (i),

$$\begin{aligned} V_2^{t=1}(1,1) - V_2^{t=1}(1,0) &\leq \mu_2 + \beta(\mathbf{V}^{t=2}(1,1) - \mathbf{V}^{t=2}(1,0)) \\ &= \frac{\partial W}{\partial a_{21}}(a_{11} = 1) \\ &< 0 \end{aligned} \tag{36}$$

where the weak inequality follows since $V_1^{t=2}(1,1) - V_1^{t=2}(1,0) \geq 0$. (36) also holds for social optimum condition (ii), since $\partial W(a_{11} = 1)/\partial a_{21} < \partial W(a_{11} = 0)/\partial a_{21} < 0$ due to Assumption 1. Under social optimum condition (iii),

$$\begin{aligned} V_1^{t=1}(1,1) - V_1^{t=1}(0,1) + V_2^{t=1}(1,1) - V_2^{t=1}(1,0) \\ \leq \frac{\partial W}{\partial a_{11}}(a_{21} = 1) + \frac{\partial W}{\partial a_{21}}(a_{11} = 1) \\ < 0 \end{aligned} \tag{37}$$

where the weak inequality follows since $V_1^{t=2}(1,1) - V_1^{t=2}(1,0)$ and $V_2^{t=2}(1,1) - V_2^{t=2}(0,1)$ are non-negative. This shows that the decentralized outcome cannot have both policymakers experimenting when the social optimum is $(a_{11} = 1, a_{21} = 0)$. Instead, decentralization must involve either one or no policymaker experimenting. It is not difficult to find parameter values where the latter case holds.³⁷

³⁷For example, when $\mu_2 \rightarrow -\infty$, policymaker 2 does not experiment in the decentralized or social optimum and has a payoff of zero. When also $\rho > 0$ (so $\sigma_1^{*2}[1,1] > \sigma_1^{*2}[1,0]$), then $\partial W(a_{21} = 1)/\partial a_{11} - V_1^{t=1}(1,0) = \beta(V_1^{t=2}(1,1) - V_1^{t=2}(1,0))$ is strictly positive. So there must exist parameter values which involve no experimentation under decentralization ($V_1^{t=1}(1,0) < 0$) when condition (i) holds.

The only remaining case is when both policymakers experiment in the social optimum, ($a_{11} = 1, a_{21} = 1$). By definition there cannot be over-experimentation though it is possible to find parameter values which imply under-experimentation.³⁸ This completes the proof. ■

A.4 Proof of Proposition 1: The proof proceeds by arguing that the centralized payoff from experimenting is non-decreasing in ρ , positive when $\rho = 1$, and negative when $\rho = 0$. Now differentiating (18) with respect to ρ when $a_1 = 1$ gives,

$$\frac{\partial \tilde{V}^{t=1}}{\partial \rho} = \beta \phi \left(\frac{\mu}{\sigma^*[1]} \right) \frac{\partial \sigma^*[1]}{\partial \rho} \tag{38}$$

where I have used $\partial \tilde{V}^{t=2}(a_1)/\partial \sigma^*[a_1] = \phi(\mu_i/\sigma^*[a_1])$. The left-hand side of (38) has the same sign as $\partial \sigma^*[1]/\partial \rho$. Now the covariance term in (17) is the off-diagonal term in (5) when α is the identity matrix and can be written as,

$$\text{Cov}[1,1] \equiv \frac{(1 - \rho^2)\sigma_1^2\sigma_2^2 + \sigma_1^2 + \sigma_2^2}{(1 - \rho^2)\sigma_1^2\sigma_2^2 + \sigma_1^2 + \sigma_2^2 + 1} \rho\sigma_1\sigma_2 \tag{39}$$

Substituting (26) and (39) into (17) and differentiating gives,

$$\begin{aligned} \frac{\partial \sigma^*[1]}{\partial \rho} = & \frac{[(\sigma_1^2 + \sigma_2^2 + 2)\rho + ((1 - \rho^2)\sigma_1^2\sigma_2^2 + \sigma_1^2 + \sigma_2^2 + 3) \\ & \times (1 - \rho^2)\sigma_1\sigma_2] \sigma_1\sigma_2 \\ & + ((1 - \rho^2)\sigma_1^2\sigma_2^2 + \sigma_1^2 + 1)\sigma_1^2 + ((1 - \rho^2)\sigma_1^2\sigma_2^2 + \sigma_2^2 + 1)\sigma_2^2}{((1 - \rho^2)\sigma_1^2\sigma_2^2 + \sigma_1^2 + \sigma_2^2 + 1)^2} \\ & \times 2\sigma_1\sigma_2 \end{aligned} \tag{40}$$

(40) is clearly positive when $\rho \geq 0$. Since the numerator of (40) is increasing in ρ when $\rho < 0$, if this derivative is ever negative it must be so at $\rho = -1$. Substituting $\rho = -1$ into (40) yields,

$$\frac{\partial \sigma^*[1]}{\partial \rho} (\rho = -1) = \frac{(\sigma_1^3 - \sigma_2^3) + (\sigma_1 - \sigma_2)}{(\sigma_1^2 + \sigma_2^2 + 1)^2} (\sigma_1 - \sigma_2) 2\sigma_1\sigma_2 \tag{41}$$

which must be non-negative. This shows that $\partial \sigma^*[1]/\partial \rho \geq 0$ and so $\partial \tilde{V}^{t=1}/\partial \rho \geq 0$ with equality only at $\rho = -1$ and $\sigma_1 = \sigma_2$.

The sure policy is non-dominant (Assumption 2). Therefore, the sure policy must give policymaker i a lower payoff than the experimental policy when $-i$ is using the sure policy (due to Assumption 1). Define $\underline{\mu}_i$ to be the smallest μ_i at which this condition just holds (conditional on the

³⁸For example, when $\rho \neq 0$ then $\partial W(a_{11} = 1)/\partial a_{21} - (V_2^{t=1}(1,1) - V_2^{t=1}(1,0)) = \beta(V_1^{t=2}(1,1) - V_1^{t=2}(1,0))$ is strictly positive. This means there are parameter values where the social optimum has both policymakers experimenting ($\partial W(a_{-i1} = 1)/\partial a_{i1} \geq 0$), but the decentralized outcome has policymaker 2 using the sure policy ($V_2^{t=1}(1,1) < V_2^{t=1}(1,0)$).

other parameter values). Since the sure policy gives a zero payoff when no one experiments here, $\underline{\mu}_1$ is characterized by,

$$\underline{\mu}_1 + \beta \left[\underline{\mu}_1 \Phi \left(\frac{\underline{\mu}_1}{\sigma_1^*[1,0]} \right) + \sigma_1^*[1,0] \phi \left(\frac{\underline{\mu}_1}{\sigma_1^*[1,0]} \right) \right] = 0 \tag{42}$$

using (8) and (9). The equation characterizing $\underline{\mu}_2$ is similar except it has $\sigma_2^*[0,1]$ rather than $\sigma_1^*[1,0]$. Notice that (42) can be re-written as,

$$\underline{c}_1 + \beta [\underline{c}_1 \Phi(\underline{c}_1) + \phi(\underline{c}_1)] = 0 \tag{43}$$

where $\underline{c}_1 \equiv \underline{\mu}_1 / \sigma_1^*[1,0]$. Again a similar equation must hold for policymaker 2 substituting $\underline{c}_2 \equiv \underline{\mu}_2 / \sigma_2^*[0,1]$ for \underline{c}_1 . Because the left-hand side of (43) is monotone increasing in \underline{c}_1 , it must be that $\underline{c}_1 = \underline{c}_2$ or,

$$\underline{\mu}_2 = \underline{\mu}_1 \frac{\sigma_2^*[0,1]}{\sigma_1^*[1,0]} \tag{44}$$

In the centralized case experimentation is optimal if it yields an expected payoff greater than zero, the return from the sure policy. Using (16) and (18) the relevant cut-off is characterized by,

$$\underline{c} + \beta [\underline{c} \Phi(\underline{c}) + \phi(\underline{c})] = 0 \tag{45}$$

where $\underline{c} \equiv \underline{\mu} / \sigma^*[1]$. Because (45) is similar to (43), it must be that $\underline{c} = \underline{c}_i$. To show that centralization involves experimenting when $\mu_i = \underline{\mu}_i$, it will be sufficient to show $\underline{\mu} / \sigma^*[1] > \underline{c} = \underline{c}_i$ where $\underline{\mu} \equiv \underline{\mu}_1 + \underline{\mu}_2$ (this is because the left-hand side of (45) is monotone increasing in \underline{c}). Using (44) and $\underline{\mu}_i < 0$ (which is necessary for non-dominance), this condition is,

$$\sigma^{*2}[1] \geq (\sigma_1^*[1,0] + \sigma_2^*[0,1])^2 \tag{46}$$

Now the right-hand side of (46) is independent of ρ (see (24)) while I showed above that the left-hand side is increasing in ρ . To prove the result, it will be sufficient to find *some* ρ for which (46) holds. When $\rho = 1$, $\text{Cov}[1,1] = \sigma_1^*[1,1]\sigma_2^*[1,1]$ and so using (17) $\sigma^{*2} = (\sigma_1^*[1,1] + \sigma_2^*[1,1])^2$. Since $\rho = 1$ also means that $\sigma_1^*[1,1] > \sigma_1^*[1,0]$ and $\sigma_2^*[1,1] > \sigma_2^*[0,1]$, (46) must hold.

Now the above result is predicated on $\mu_i = \underline{\mu}_i$ whereas larger μ_i are also compatible with decentralized mixing. It is not difficult to show that $\partial \tilde{V}^{t=1} / \partial \mu_i > 0$, so the above result shows that when there is decentralized mixing that centralization must involve experimenting when $\rho = 1$. Finally, it is possible to show that centralized experimentation cannot occur when $\rho = 0$. Now when $\rho = 0$ it must be that $\mu_i = \underline{\mu}_i$ because otherwise Assumption 2 fails (experimenting is dominant). To see this notice that the numerator of (12), which is identical to (42), must be 0 because the denominator is also 0 when $\rho = 0$. So to show

there cannot be centralized experimenting, it is sufficient to show (46) fails. When $\rho = 0$, $\text{Cov}[1,1] = 0$ and so $\sigma^{*2}[1] = \sigma_1^{*2}[1,1] + \sigma_2^{*2}[1,1]$. Since $\rho = 0$ also means that $\sigma_1^*[1,1] = \sigma_1^*[1,0]$ and $\sigma_2^*[1,1] = \sigma_2^*[0,1]$, (46) fails. Because $\partial \tilde{V}^{t=1} / \partial \rho \geq 0$ with equality only at $\rho = -1$ and $\sigma_1 = \sigma_2$, this proves the main claim: $\exists! \rho^* > 0 : a_1 = 1 \forall \rho \geq \rho^*$. ■

A.5 *Proof of Proposition 2:* The proof proceeds by finding the conditions under which centralization involves experimentation.³⁹ Under centralized experimentation the inverse of the variance matrix in (19) is,

$$\Sigma_N^{-1} = \Delta_N^{-1} \begin{pmatrix} (N-2)\rho + 1 & -\rho & -\rho & \dots \\ -\rho & (N-2)\rho + 1 & -\rho & \dots \\ & & \ddots & \\ -\rho & -\rho & \dots & (N-2)\rho + 1 \end{pmatrix} \quad (47)$$

where $\Delta_N \equiv ((N-2)\rho + 1 - (N-1)\rho^2)\sigma^2$. Under centralized experimentation, α is the identity matrix and (5) can be re-written as $\Sigma_N^{-1} \text{Cov}_N(\boldsymbol{\mu}^*) + \text{Cov}_N(\boldsymbol{\mu}^*) = \Sigma_N^{-1}$. Using (47) this equation can be solved to yield,

$$\text{Cov}_N(\boldsymbol{\mu}^*) = \begin{pmatrix} \text{Var}_N & \text{Cov}_N & \text{Cov}_N & \dots \\ \text{Cov}_N & \text{Var}_N & \text{Cov}_N & \dots \\ & & \ddots & \\ \text{Cov}_N & \text{Cov}_N & \dots & \text{Var}_N \end{pmatrix} \quad (48)$$

where,

$$\text{Var}_N \equiv \frac{1 + \sigma^2 + (N-2)\rho\sigma^2 - (N-1)\rho^2(\sigma^2 - 1)}{(1 + (1-\rho)\sigma^2)(1 + (1 + (N-1)\rho)\sigma^2)} \sigma^4 \quad (49)$$

and,

$$\text{Cov}_N \equiv \frac{2 + \sigma^2 - (N-1)\rho^2\sigma^2 + (N-2)\rho(1 + \sigma^2)}{(1 + (1-\rho)\sigma^2)(1 + (1 + (N-1)\rho)\sigma^2)} \rho\sigma^4 \quad (50)$$

Following the proof in section A.4, initially presume that the experimental means are at the lowest level compatible with non-dominance, $\mu_i = \underline{\mu}_i$. The condition characterizing centralized experimentation is $(\sum_i \mu_i / \sigma^*[N]) > \underline{c}$. Here \underline{c} is determined by (45), and $\sigma^{*2}[N]$ is the variance of the expected second-period experimental mean,

$$\sigma^{*2}[N] \equiv N\text{Var}_N + N(N-1)\text{Cov}_N \quad (51)$$

³⁹In an earlier version of this paper I show that decentralization involves a unique number of first-period experiments when pure actions are used.

Notice that (51) is simply a generalization of (17). The cut-off condition for experimentation may be written as,

$$\sigma^{*2}[N] \geq (N\sigma^*[1,0])^2 \tag{52}$$

using $\sigma_i^2 = \sigma^2 \forall i$, $\underline{c} = \underline{c}_1 \equiv \mu_1/\sigma^*[1,0]$, and $\mu_1 < 0$ (which is necessary for non-dominance). A comparison of (24) with (49)–(51) yields,

$$\begin{aligned} &\sigma^{*2}[N] - (N\sigma^*[1,0])^2 \\ &= \frac{-1 - \sigma^2 + (N-1)\rho^2(1 + \sigma^2) + \rho(2 - (N-2)\sigma^2)}{(1 + \sigma^2)(1 + (1 + (N-1)\rho)\sigma^2)} N(N-1)\sigma^4 \end{aligned} \tag{53}$$

which can be positive or negative. Treating the numerator of (53) as a quadratic in ρ and setting it equal to zero gives the roots,

$$\rho^*(N), \rho^{**}(N) \equiv \frac{(-2 + (N-2)\sigma^2) \pm ((4 + 4\sigma^2 + N\sigma^4)N)^{0.5}}{2(N-1)(1 + \sigma^2)} \tag{54}$$

The positive root $\rho^*(N)$ is greater than zero, less than one, and is decreasing in N . The negative root $\rho^{**}(N)$ is less than zero. The denominator of (53) is negative when ρ is less than,

$$\tilde{\rho}(N) \equiv \frac{-(1 + \sigma^2)}{\sigma^2(N-1)} < \rho^{**}(N) \tag{55}$$

Now the condition for centralized experimenting is that (53) is positive. Given the above results, this only occurs when: (i) $\rho > \rho^*(N)$; or (ii) $\rho \in [\tilde{\rho}(N), \rho^{**}(N)]$. However, (ii) cannot occur because $\rho^{**}(N) < -(N-1)^{-1}$ and $\rho \geq -(N-1)^{-1}$ is necessary for Σ_N to be positive semidefinite (see note 27).

Raising μ_i above $\underline{\mu}_i$ increases the benefit of centralized experimentation and so reduces the cut-off $\rho^*(N)$. However, the cut-off must remain positive so long as experimentation is non-dominant for all regions. The comparative statics with respect to N discussed above still hold here since the higher μ_i values are not a function of N . ■

A.6 Proof of Proposition 3: The proof proceeds by finding conditions under which the two decentralized policymakers use different experimental policies. The first step is to show that it may be optimal for a policymaker to use experimental policy j rather than $-j$ when his neighbor is using $-j$. This requires showing $V_1^{t=2}(1,0;0,1) > V_1^{t=2}(0,0;1,1)$. $V_1^{t=2}(0,0;1,1)$ can be written as in (8), since this is simply the truncated expected value of an experimental policy. $V_1^{t=2}(1,0;0,1)$ is the truncated expected maximum of two normally distributed variables. Applying the approach in David (1981) to (20), this term may be written as,

$$\begin{aligned}
 V_1^{t=2}(1,0;0,1) &= \int_{\frac{-\mu_1}{\sigma_1^*[1,0]}}^{\infty} (\sigma_1^*[1,0]y + \mu_1)\Phi\left(\frac{\sigma_1^*[1,0]y}{\sigma_1^*[0,1]}\right)\phi(y) dy \\
 &+ \int_{\frac{-\mu_1}{\sigma_1^*[0,1]}}^{\infty} (\sigma_1^*[0,1]y + \mu_1)\Phi\left(\frac{\sigma_1^*[0,1]y}{\sigma_1^*[1,0]}\right)\phi(y) dy \quad (56)
 \end{aligned}$$

The $\sigma_1^*[\cdot]$ terms are defined in section A.1 in the case where $\sigma_1 = \sigma_2 \equiv \sigma$, and $\mu_1 \equiv \mu$. While the integrals in (56) do not have a general closed form solution, expressions can be written for special cases. When $\rho = 0$ there is no information spillover and so $V_1^{t=2}(1,0;0,1) = V_1^{t=2}(0,0;1,1)$. Numerical solutions suggest that if experimenting with different policies is optimal, it must be so at $\rho \rightarrow \pm 1$.⁴⁰ When $\rho = \pm 1$, $\sigma_1^*[1,0] = \sigma_1^*[0,1]$ and (56) reduces to,

$$\begin{aligned}
 V_1^{t=2}(1,0;0,1) &= 2\sigma_1^*[1,0]\left(1 - \Phi\left(\frac{\mu_1}{\sigma_1^*[1,0]}\right)\right)\phi\left(\frac{\mu_1}{\sigma_1^*[1,0]}\right) \\
 &+ \sqrt{2}\sigma_1^*[1,0]\Phi\left(\frac{\sqrt{2}\mu_1}{\sigma_1^*[1,0]}\right)\phi(0) \\
 &+ \mu_1\left(1 - \left(1 - \Phi\left(\frac{\mu_1}{\sigma_1^*[1,0]}\right)\right)^2\right) \quad (57)
 \end{aligned}$$

This can be used to show $V_1^{t=2}(1,0;0,1) > V_1^{t=2}(0,0;1,1)$ is possible. For example, when $\rho = \pm 1$ and $\mu_1 \rightarrow 0^-$ then a comparison of (8) and (57) shows the inequality holds (because $\sqrt{2}\sigma_1^*[1,0] > \sigma_1^*[1,1]$ when $\rho = \pm 1$). By continuity this inequality (and those below) also hold for ρ values near ± 1 and μ_1 values near 0. They also hold for policymaker 2 since $\mu_{ij} = \mu$ and $\sigma_{ij} = \sigma$.

The next step is to show that a policymaker may use experimental policy j rather than the sure policy when his neighbor uses experimental policy $-j$. This requires showing $\mu_1 + \beta V_1^{t=2}(1,0;0,1) > \beta V_1^{t=2}(0,0;0,1)$. This condition holds for a variety of parameter values such as the ones used in the last paragraph, $\rho = \pm 1$, $\mu_1 \rightarrow 0^-$.

When the inequalities in the last two paragraphs hold, there is a decentralized equilibrium where the policymakers use different experimental policies.⁴¹ This is because the above results show that the optimal response to a neighbor's experiment is to use the other experimental policy. Now the previous paragraphs show that the policymakers' second-period return is maximized when each experi-

⁴⁰It is not possible to formally prove this point since $\partial V_1^{t=2}(1,0;0,1)/\partial \rho \cong \partial V_1^{t=2}(0,0;1,1)/\partial \rho$ and $\partial^2 V_1^{t=2}(1,0;0,1)/\partial \rho^2 \cong \partial^2 V_1^{t=2}(0,0;1,1)/\partial \rho^2$.

⁴¹For such parameter values there may also be an equilibrium where neither policymaker experiments. This possibility is ruled out with another extension of Assumption 1, $V_1^{t=2}(0,1;0,0) + V_1^{t=2}(1,0;0,0) - V_1^{t=2}(1,0;0,1) - V_1^{t=2}(0,0;0,0) > 0$.

ments with a different policy. This means aggregate experimentation is higher under decentralization than centralization, since the latter can only involve a one policy experiment. Notice also that decentralization can have both experimental policies being used even when centralization has no experimentation at all. When $\rho = -1$, $\text{Cov}[1,1] = -\sigma^{*2}[1,1]$ and so using (17) $\sigma^{*2}[1] = 0$. This means that centralization will not involve experimentation so long as $\mu < 0$ while the above analysis shows that decentralization can involve each policymaker using a different experimental policy. ■

References

- BALA, V., and S. GOYAL (1998) Learning from neighbors, *Review of Economics Studies* **65**, 595–622.
- BERGEMANN, D., and J. VALIMAKI (1997) Market diffusion with two-sided learning, *RAND Journal of Economics* **28**, 773–795.
- BERRY, F. S. (1994) Sizing up state policy innovation research, *Policy Studies Journal* **22**, 442–456.
- BERRY, F. S., and W. D. BERRY (1990) State lottery adoptions as policy innovations: an event history analysis, *American Political Science Review* **84**, 395–415.
- BESLEY, T., and A. CASE (1995) Incumbent behavior: vote-seeking, tax-setting, and yardstick competition, *American Economic Review* **85**, 25–45.
- BESLEY, T., and S. COATE (1999) “Centralized versus decentralized provision of local public goods: a political economy analysis,” NBER working paper no. 7084.
- BINGHAM, R. D. (1976) *The Adoption of Innovation by Local Governments*. Lexington, MA: Lexington Books.
- BLANK, R. M. (1997) Policy watch: the 1996 welfare reform, *Journal of Economic Perspectives* **11:1**, 169–177.
- BOLTON, P., and C. HARRIS (1999) Strategic experimentation, *Econometrica* **67**, 349–374.
- CANON, B. C., and L. BAUM (1981) Patterns of adoption of tort law innovations: an application of diffusion theory to judicial doctrines, *American Political Science Review* **75**, 975–987.
- CHANG, M.-H., and J. HARRINGTON (forthcoming) Centralization vs. decentralization in a multi-unit organization: a computational model of a retail chain as a multi-agent adaptive system, *Management Science*.
- CREMER, J., and T. R. PALFREY (1998) “Federal mandates by popular demand,” unpublished working paper.
- DAVID, H. A. (1981) *Order Statistics*. New York: John Wiley & Sons.
- DeGROOT, M. H. (1970) *Optimal Statistical Decisions*. New York: McGraw-Hill.
- DIXIT, A., and J. LONDREGAN (1998) Fiscal federalism and redistributive politics, *Journal of Public Economics* **68**, 153–180.

- FREEMAN, P. K. (1985) Interstate communication among state legislators regarding energy policy innovation, *Publius* **15:4**, 99–112.
- GLICK, H. (1981) Innovation in state judicial administration: effects on court management and organization, *American Politics Quarterly* **9**, 49–69.
- INMAN, R. P., and D. L. RUBINFELD (1997) Rethinking federalism, *Journal of Economic Perspectives* **11:4**, 43–64.
- JOHNSON, N. L., and S. KOTZ (1970) *Continuous Univariate Distributions-1*. New York: Houghton Mifflin.
- KOLLMAN, K., J. H. MILLER, and S. E. PAGE (1997) “Searching for policy solutions in a federal system,” unpublished working paper.
- McCOY, E. C. (1940) Patterns of diffusion in the United States, *American Sociological Review* **5**, 219–227.
- MINTROM, M. (1997) Policy entrepreneurs and the diffusion of innovation, *American Journal of Political Science* **41**, 738–770.
- NICE, D. C. (1994) *Policy Innovation in State Government*. Ames: Iowa State University Press.
- OATES, W. E. (1999) An essay on fiscal federalism, *Journal of Economic Literature* **37**, 1120–1149.
- PERSSON, T., and G. TABELLINI (1996a) Federal fiscal constitutions: risk sharing and moral hazard, *Econometrica* **64**, 623–646.
- PERSSON, T., and G. TABELLINI (1996b) Federal fiscal constitutions: risk sharing and redistribution, *Journal of Political Economy* **104**, 979–1009.
- REGENS, J. L. (1980) State policy responses to the energy issue, *Social Science Quarterly* **61**, 44–57.
- ROSE-ACKERMAN, S. (1980) Risk taking and reelection: does federalism promote innovation? *The Journal of Legal Studies* **9**, 593–616.
- ROSE-ACKERMAN, S. (1981) Does federalism matter? Political choice in a federal republic, *Journal of Political Economy* **89**, 152–165.
- ROSEN, H. S. (1999) *Public Finance*. Boston: McGraw-Hill.
- SIGELMAN, L., et al. (1981) Social service innovation in the American states, *Social Science Quarterly* **62**, 503–515.
- SMITH, L., and P. SORENSEN (forthcoming) Pathological outcomes of observational learning, *Econometrica*.
- STRUMPF, K. S., and F. OBERHOLZER-GEE (forthcoming) Endogenous policy decentralization: testing the central tenet of economic federalism, *Journal of Political Economy*.
- WALKER, J. L. (1971) Innovation in state politics, in H. Jacob and K. Vines, eds., *Politics in the American States*. Boston: Little, Brown.
- WISEMAN, M. (1996) State strategies for welfare reform: the Wisconsin story, *Journal of Policy Analysis and Management* **15**, 515–546.